


05/01/98



JCS29 U.S. PTO

Please type a plus sign (+) inside this box → ☐

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL (Only for new nonprovisional applications under 37 CFR 1.53(b))	Attorney Docket No.	15886-122	Total Pages	33
	First Named Inventor or Application Identifier			
	Glenn William Connery; Title: Receive Processing With Network Protocol Bypass			
	Express Mail Label No.	EM033395650 US		

APPLICATION ELEMENTS See MPEP chapter 600 concerning utility patent application contents.		ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231	
1. <input checked="" type="checkbox"/> Fee Transmittal Form (Submit an original, and a duplicate for fee processing) 2. <input checked="" type="checkbox"/> Specification [Total Pages <u>18</u>] (Preferred arrangement set forth below) - Descriptive title of the Invention - Cross References to Related Applications - Statement Regarding Fed sponsored R&D - Reference to Microfiche Appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings - Detailed Description - Claim(s) - Abstract of the Disclosure 3. <input checked="" type="checkbox"/> Drawing(s) (37 CFR 1.152) [Total Sheets <u>4</u>] 4. <input checked="" type="checkbox"/> Oath or Declaration [Total Pages <u>4</u>] a. <input checked="" type="checkbox"/> Newly executed (original or copy) b. <input type="checkbox"/> Copy from a prior application (37 CFR 1.63(d)) (for continuation/divisional with Box 17 completed) [Note Box 5 below] i. <input checked="" type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b). 5. <input type="checkbox"/> Incorporation By Reference (useable if Box 4b is checked) The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.		6. <input type="checkbox"/> Microfiche Computer Program (Appendix) 7. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy (identical to computer copy) c. <input type="checkbox"/> Statement verifying identity of above copies 8. <input checked="" type="checkbox"/> Assignment Papers (cover sheet & documents(s)) 9. <input type="checkbox"/> 37 CFR 3.73(b) Statement (when there is an assignee) <input type="checkbox"/> Power of Attorney 10. <input type="checkbox"/> English Translation Document (if applicable) 11. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations 12. <input type="checkbox"/> Preliminary Amendment 13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) (Should be specifically itemized) 14. <input type="checkbox"/> Small Entity Statement(s) <input type="checkbox"/> Status still proper and desired 15. <input type="checkbox"/> Certified Copy of Priority Document(s) (if foreign priority is claimed) 16. <input checked="" type="checkbox"/> Other: Power of Attorney By Assignee To Exclusion Of Inventor Under 37 C.F.R. §3.71 With Revocation Of Prior Powers	

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No. /

16. CORRESPONDING ADDRESS					
<input type="checkbox"/> Customer Number of Bar Code Label (Insert Customer No. or Attach bar code label here)					
or <input checked="" type="checkbox"/> Correspondence address below					
NAME	Mark A. Haynes Reg. No. 30,846				
ADDRESS	WILSON SONSINI GOODRICH & ROSATI 650 Page Mill Road				
CITY	Palo Alto	STATE	California	ZIP CODE	94304-1050
COUNTRY	USA	TELEPHONE	(650) 493-9300	FAX	(650) 845-5000

SUBMITTED BY

Typed or

Printed Name Mark A. HaynesSignature Mark A. HaynesReg. Number 30,846Date 1 May 98

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

FEE TRANSMITTALNote: Effective October 1, 1997.
Patent fees are subject to annual revision.**Complete if Known**

TOTAL AMOUNT OF PAYMENT (\$)	830.00	Application Number	Not Yet Assigned
		Filing Date	Herewith
		First Named Inventor	Glenn William Connerly, Title: Receive Processing With Network Protocol Bypass
		Group Art Unit	Not Yet Assigned
		Examiner Name	Not Yet Assigned
		Attorney Docket Number	15886-122

METHOD OF PAYMENT (check one)

1. ☒ The Commissioner is hereby authorized to charge indicated fees and credit any over payments to:

Deposit Account Number: 23-2415 (15886-122)

Deposit Account Name: Wilson Sonsini Goodrich & Rosati

☒ Charge Any Additional Fee Required Under 37 CFR 1.16 and 1.17 ☐ Charge the Issue Fee Set in 37 CFR 1.18 at the Mailing of the Notice of Allowance

2. ☐ Payment Enclosed: ☐ Check ☐ Money Order ☐ Other

FEE CALCULATION**1. FILING FEE**

Larg e	Entity	Small	Entity	Fee Code	Fee (\$)	Fee Code	Fee (\$)	Fee Description	Fee Paid
101	790	201	395					Utility filing fee	790.00
106	330	206	165					Design filing fee	
107	540	207	270					Plant filing fee	
108	790	208	395					Reissue filing fee	
114	150	214	75					Provisional filing fee	
SUBTOTAL (1) (\$)									790.00

2. CLAIMS

Total Claims	Extra	Fee from below	Fee Paid
20 - 20 =	0	22	0.00
Independent Claims 3 - 3 =	0	82	0.00
Multiple Dependent Claims	X	270	

Larg e	Entity	Small	Entity	Fee Code	Fee (\$)	Fee Code	Fee (\$)	Fee Description	Fee Paid
103	22	203	11					Claims in excess of 20	
102	82	202	41					Independent claims in excess of 3	
104	270	204	135					Multiple dependent claim	
109	82	209	41					Reissue independent claims over original patent	
110	22	210	11					Reissue claims in excess of 20 and over original patent	
SUBTOTAL (2) (\$)									0.00

FEE CALCULATION (continued)**3. ADDITIONAL FEES**

Larg e	Entity	Small	Entity	Fee Code	Fee (\$)	Fee Code	Fee (\$)	Fee Description	Fee Paid
105	130	205	65					Surcharge - late filing fee or oath	
127	50	227	25					Surcharge - late provisional filing fee or cover sheet	
139	130	139	130					Non-English specification	
147	2,520	147	2,520					For filing a request for reexamination	
112	920*	112	920*					Requesting publication of SIR prior to Examiner action	
113	1,840*	113	1,840*					Requesting publication of SIR after Examiner action	
115	110	215	55					Extension for reply within first month	
116	400	216	200					Extension for reply within second month	
117	950	217	475					Extension for reply within third month	
118	1,510	218	755					Extension for reply within fourth month	
128	2,060	228	1,030					Extension for reply within fifth month	
119	310	219	155					Notice of Appeal	
120	310	220	155					Filing a brief in support of an appeal	
121	270	221	135					Request for oral hearing	
138	1,510	138	1,510					Petition to institute a public use proceeding	
140	110	240	55					Petition to revive - unavoidable	
141	1,320	241	660					Petition to revive - unintentional	
142	1,320	242	660					Utility issue fee (or reissue)	
143	450	243	225					Design issue fee	
144	670	244	335					Patent issue fee	
122	130	122	130					Petitions to the Commissioner	
123	50	123	50					Petitions related to provisional applications	
126	240	126	240					Submission of Information Disclosure Stmt	
581	40	581	40					Recording each patent assignment per property (times number of properties)	40.00
146	790	246	395					Filing a submission after final rejection (37 CFR 1.129(a))	
149	790	249	395					For each additional invention to be examined (37 CFR 1.129(b))	
Other fee (specify)									
Other fee (specify)									
SUBTOTAL (3) (\$)									40.00

* Reduced by Basic Filing Fee Paid

SUBTOTAL (3) (\$) 40.00**SUBMITTED BY**

Typed or Printed Name	Mark A. Haynes	Complete (if applicable)	
Signature	<i>Mark A. Haynes</i>	Reg. Number	30,846
Date	1/16/97	Deposit Account	23-2415
		User ID	(15886-122)

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Washington, DC 20231.

RECEIVE PROCESSING WITH NETWORK PROTOCOL BYPASS

Inventors: Glenn William Connery
Gary Jaszewski
Richard Reid

5

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to processing of data in communication networks, and more particularly to the process of receiving a plurality of packets of data which relate to a common block of data, and efficiently providing such data to an application.

10

Description of Related Art

Network communications are often described with respect to layers of network protocols. According to a standard description, the layers include the physical layer, the datalink layer, the network layer (also called routing layer), the transport layer, and the application layer. Thus modern communication standards, such as the Transport Control Protocol TCP, the Internet Protocol IP, and IEEE 802 standards, can be understood as organizing the tasks necessary for data communications into layers. There are a variety of types of protocols that are executed at each layer according to this model. The particular protocols utilized at each layer are mixed and matched in order to provide so called protocol stacks or protocol suites for operation of a given communication channel.

15

20

25

The protocol stacks typically operate in a host system which includes a network adapter comprised of hardware that provides a physical connection to a network medium, and software instructions referred to as medium access control MAC drivers for managing the communication between the adapter hardware and the protocol stack in the host system. The adapter generally

includes circuitry and connectors for communication over a communication medium, and translates the data to and from the digital form used by the protocol stack and the MAC driver, and a form that may be transmitted over the communication medium.

5 Generally according to this model, processes at the application layer, including applications and file systems, rely on the lower layers of the communication protocol stack for transferring the data between stations in the network. The application layer requests services from the protocol stack which includes transport layer, network layer and datalink layer processes distributed
10 between the MAC driver and other components of the stack. In a similar way, data which is received across the network is passed up the protocol stack to the application layer at which actual work on the data involved is accomplished.

In current implementations, received packets are generally moved sequentially into host buffers allocated by the MAC driver for the adapter, as
15 they arrive. These buffers are then provided to the host protocol stack, which generally copies them once or twice to internal buffers of its own before the payload data finally gets copied to the application or the file system buffer. This sequential passing of the data up the protocol stack is required so that the processes in the particular protocol suite are able to individually handle the
20 tasks necessary according to the protocol at each layer. However, these multiple copies of the data hurt performance of the system. In particular, the CPU of the computer is used for each copy of the packet, and a significant load is placed on the memory subsystem in the computer. With technologies like gigabit Ethernet, and other technology in which the data rates of the physical layer of
25 the network is increasing, these copy operations may become an important limiting factor in improving performance of personal computer architectures to levels approaching the capability of the networks to which they are connected.

Accordingly, it is desirable to provide techniques which avoid one or more of these copies of the packets as they pass up the protocol stacks. By

eliminating multiple copies of the packet, the raw performance of the receiving end station can be increased, and the scalability of the receive process can be improved.

5

SUMMARY OF THE INVENTION

According to the present invention, an adapter is provided with intelligence that allows it to separate the header parts of a packet being received from the payload it carries, and in most cases move the payload directly into a destination buffer at a higher layer, such as the application layer. Thus reducing the number of times that the payload of a communication must be copied by the host system.

Accordingly, the invention can be characterized as a method for transferring data on a network from the data source to an application executing in an end station. The application operates according to a multi-layer network protocol which includes a process for generating packet control data (e.g. headers) for packets according to the multi-layer network protocol. Packets are received at the network interface in a sequence carrying respective data payloads from the data source. Upon receiving a packet, the control data of the packet is read in the network interface, and if the packet belongs to a flow specification subject of the bypass, the data payload of the packet is transferred to a buffer assigned by a layer higher in the stack, preferably by the application or file system, bypassing one or more intermediate buffers of the protocol stack.

Typically, to initiate the process of receiving a plurality of packets which make up a block of data for a particular application, the process involves establishing a connection between the end station and the source of data, such as a file server on a network, for example according to the TCP/IP protocol suite. A request is transmitted from the application through the network interface which asks for transfer of the data from the data source. The request and the protocol suite provide a flow specification to identify the block of data and an

BRIEF DESCRIPTION OF THE FIGURES

Fig. 1 is a simplified diagram of an end station in a network including receive bypass support according to the present invention.

Fig. 2 is a simplified diagram of a network interface card supporting the receive bypass processes of the present invention.

Fig. 3 is a heuristic diagram of a protocol stack including the read bypass process of the present invention.

Fig. 4 is a diagram illustrating the packet structure and processing according to the read bypass operation of the present invention.

Fig. 5 is a flow diagram of the process executed for setting up and receiving packets in response to a READ RAW request.

DETAILED DESCRIPTION

A detailed description of the present invention is provided with respect to Figs. 1-5, in which Figs. 1 and 2 illustrate a hardware system environment.

Fig. 1 shows a data processing system 10 which includes a host central processing unit 11, host memory 12, host input/output 13, such as keyboards, displays, printers, a pointing device and the like. The system also includes program memory 14 (usually part of the host memory block) and a network interface card 15. All of these elements are interconnected by a host system bus 16. The network interface card 15 provides for connection to a network medium as indicated at line 17.

Fig. 1 is a simplified diagram of a computer such as a personal computer or workstation. The actual architecture of such system is quite varied. This system for one example corresponds to a personal computer based on the Intel microprocessor running a Microsoft Windows operating system. Other combinations of processor and operating system are also suitable.

According to the present invention, the program memory includes a TCP/IP protocol stack with a receive bypass mode according to the present

invention. A MAC driver is also included in the program memory which supports the receive bypass mode. Other programs are also stored in program memory to suit the needs of a particular system. The network interface card 15 includes resources to manage TCP/IP processing and bypass according to the present invention.

Fig. 2 provides a simplified block diagram of a network interface card 15 capable of supporting the present invention. The network interface card 15 includes a bus interface 20 coupled to the host bus 16. A memory composed of random access memory RAM 21 is included on the card 15. Also, a medium access control unit 22 is coupled to the card which is coupled to the network interface 17. In the path from the host bus interface 20 to the RAM 21 includes appropriate buffering 25 and a DMA engine 26 in order to offload processing from the host system for transferring data packets into the RAM 21. Also the data path from the RAM 21 to the medium access control unit 22 includes appropriate data path and buffering logic 27 to support efficient transmission and reception of packets. A DMA engine 28 is also included on this path to provide for efficient transferring of data between the network medium 17 and the RAM 21. Also included on the card 15 is a central processing unit 30 having a program memory 31. The CPU 30 is coupled to the host bus 16 and to the RAM 21 on line 32. Also the CPU 30 generates control signals represented by the arrow 33 for controlling other elements in the network interface card 15. According to this embodiment, TCP/IP checksum logic 34 is coupled to the data path and the buffering logic 27 in the path from the RAM 21 to the network medium 17. The program memory for the CPU 30 includes transmit, receive, TCP/IP receive bypass control and other processes which manage the operation of the smart adapter card.

The block diagram illustrated in Fig. 2 provides a simplified overview of the functional units in a network interface according to the present invention. A variety of other architectures could be implemented to achieve similar functions.

For example, DMA engines 26, 28 are not strictly required here. State machines hand-shaking with each other, or other data processing resources could move data from one block to the next.

In one embodiment, all of these elements are implemented on a single integrated circuit. In another embodiment, all elements of the network interface card except for the RAM 21 are implemented on a single integrated circuit. Other embodiments include discreet components for all of the major functional blocks of the network interface card.

Fig. 3 is a simplified diagram of the network protocol layers implemented according to the present invention. The protocol layers in this example include a data application 48 such as an application layer program, coupled by path 49 to a file system 50 running CIFS in an end station. The file system 50 is coupled by a path 51 to the TCP/IP stack 52. The TCP/IP stack 52 is coupled by path 53 to the MAC driver 54. The MAC driver is coupled by path 55 to the smart network interface card 56. The smart network interface card 56 is coupled by path 57 to the network medium.

In alternative systems, application layer processes (like process 48) issue read requests through an application program interface API like WINSOCK for Windows platforms, rather than a file system (like CIFS).

According to the example in Fig. 3, the application layer process 48 (such as Lotus Notes or others), issues a request 59 to read a file of data across the network from a data source such as by the equivalent of fget (buffer length), in the "C" programming language. The file system 50 sets up a bypass, then formats the READ RAW SMB request 60 and passes it down. Such a request corresponds for example to the READ RAW SMB command found in the CIFS protocol. Of course, analogous commands and other application layer processes could be utilized. The READ RAW request is passed down the protocol stack to the driver and transferred across the network to the file server or other data

source which is the source of the file to be retrieved. The file server then sends a sequence of packets according to the SMB READ RAW response protocol.

5 The TCP/IP stack according to the present invention makes a call 61 to the MAC driver, passing to it the identification of a target buffer assigned by the application layer process, or by another process above the network layer. This identification may be an array of physical or virtual addresses, such as an address and a length, which can be utilized by the MAC driver to copy the data directly from the memory on the smart network interface card into the target buffer. Also, the call made by the TCP/IP stack sends down to the MAC driver 10 the flow specification established for this read beforehand, including the source and destination IP addresses, and the source and destination port numbers (i.e. sockets). Also the flow specification includes a SEQorigin and a SMBfirst flag in this example. In turn, the MAC driver provides the request and associated control information by a transfer 62 to the smart adapter 56. The target buffer 15 specifies where the payload should be stored when possible by the network interface card. The flow specification specifies how to identify packets that are part of this session. The SEQorigin specifies the sequence number of the first byte of the payload (excluding any SMB header) that should be stored in the target buffer. The SMBfirst flag tells the driver whether the first packet with 20 that sequence SEQorigin will have a SMB header following its TCP header. This information is used so that the control data can be cut from the first packet in the plurality of packets which are received in response to the read request.

Using this information, when the adapter receives packets in the session, it puts the payload data directly (line 63) into the target buffer. There is no 25 guarantee that the adapter will always do this; however it will be done most of the time when a target buffer has been supplied. Such received packets are passed up to the protocol one at a time as always, using the same interfacing data structures as always. The only difference is that the packet will be split across two fragments. The header (Ethernet, IP, TCP and possible SMB

headers) will be the first fragment on line 64, and will occupy a driver allocated buffer as always for use in identifying the packet, and for protocol maintenance functions. However, the payload will occupy a second fragment on line 63 which has been copied to the offset within the target buffer determined by the SEQ parameter in the header. The driver figures out the offset in the target buffer from the SEQ number in the header of the incoming packet. The offset into the target buffer is simply determined by the SEQ in the packet less the SEQorigin parameter which is provided with the READ RAW request. According to one implementation, if the packets which are responsive to the read request come to the adapter out of order, they will not be redirected to the target buffer until the first packet with the SMB header carried in is received. At that point, the size of the SMB header is easily calculated, and the amount to adjust the calculation for directly loading into the target buffer can be readily determined.

This approach allows both data that is targeted at the file system cache buffers to skip one or more copies or application layer reads. However, it need not be used to its full potential to be worthwhile. Even if one copy operation can be skipped the invention might be useful. Thus, according to an alternative embodiment the target buffers are assigned at the TCP layer rather than at the application layer. The TCP buffer target address is passed down to the MAC driver, allowing the packet to skip the copy at the IP layer.

Other protocols can be handled as well, such as the IPX and the NetBEUI/DLC protocols. Although these protocols do not use SEQ numbers which act as byte counters, but rather use packet numbers as part of the flow specification provided to the adapter. Thus, the problem of calculating the offset into the target buffer is complicated. However, if all packets in the read request are constrained to the same size, except for the last packet, the target buffer offset can be easily determined. Alternatively, the bypass might only be performed on packets received in order. Additional calculations on the smart

network adapter card can also provide the memory offsets. Similarly file system protocols other than CIFS are possible, such as FTP or NFS.

According to some embodiments of the present invention, the network adapter has resources which enable it to determine that a received packet has good data, such as checksum checking logic and the like, before transferring it up to the target buffer. The smart network adapter has the capability to maintain a record of the parts of the target buffer which have been filled in with good packets and not overwrite them. This addresses a feature of the TCP/IP protocol by which there is no guarantee that payload bytes will arrive in order, or that if they are retransmitted that they will be retransmitted in the same size chunks. Thus, the file server might send part of a previous packet along with some new data. The adapter in this embodiment is capable of properly handling this condition because of its record of the good data already stored in the target buffer.

Fig. 4 illustrates the packet structure and the process of the present invention as the packets are received for a response to a READ RAW SMB request according to the CIFS protocol. As can be seen in Fig. 4, there is a first packet in the sequence generally 100 which carries a header fragment 101 and a payload fragment 102. The header fragment includes in this example an Ethernet header 103, an Internet Protocol header 104, a TCP header 105, and a SMB header 106. The data payload 107 makes up the payload fragment 102. For a standard Ethernet style packet, the payload will have about 1.5 kilobytes of data. According to the present invention, the smart network adapter receives the packet 100 and ensures that it is a good packet by performing the IP and TCP checksum processes. Also, the header fragment 101 is cut off of the packet 100 and transferred to a control header buffer 110 which is stored in driver managed memory. The payload fragment 102 is transferred to the target buffer 111 in the application managed memory according to the target buffer addresses stored in the network interface card.

According to the READ RAW SMB process, subsequent packets in the sequence, including packet 120 and packet 130 do not carry SMB headers like the header 106 in the first packet 100. The subsequent packets are received by the smart network interface card and the TCP/IP flow specification is used to identify them as part of the READ RAW response. Once they are identified by the network interface card as part of the response, they are uploaded into the target buffer 111 at an offset determined by the SEQ parameter in the TCP header. The header fragments of the packets 120 and 130 can be passed up the protocol stack to ensure that the protocol stack is properly apprised of the sequence of packets being received.

Normally a received packet is placed in a buffer allocated by the driver and then passed up to the protocol stack in a data structure that consists of one or more fragments identified by respective pointers and lengths. The total sum of the fragments in order makes up the entire packet. Often, the packet is passed up in one piece, and there is not a second fragment. But in some cases, such as in a transmit loop back, the packet is divided into several fragments which are passed back up to the protocol stack. Thus the protocol stack is normally configured to handle packets which are passed up in several fragments. According to the present invention, the packet is divided into two fragments, including the header which is placed in the buffer allocated by the driver identified by a pointer to the buffer location, and the payload which is placed in the target buffer and identified by a pointer to the target buffer. The identifiers of the two fragments are passed up the protocol stack by making a call to the receive function in the next layer, and passing the fragment identifiers up with the call. This allows the packet to be processed normally through the protocol stack. At the application layer, or at the layer of the target buffer, the protocol or the application would be modified according to one implementation of the invention to compare the address of the fragment with the address of the buffer into which this layer of the stack intends to copy the fragment. If these

two addresses match, then the copy is not executed. The copy would not be necessary because the adapter had already copied the payload into the target buffer. Thus, very little modification of the protocol stack is necessary in order to execute the present invention.

5 Fig. 5 illustrates the process executed by the network interface card and MAC driver according to the present invention. In the network adapter the bypass is set up by storing a flow specification for the block of data that will be subject of the receive bypass. The flow specification includes for example a sequence number range, socket numbers, and source and destination addresses,
10 of packets which will be part of the flow, and a target buffer to which the payload is to be written (block 200). This data is provided to the network adapter from a higher layer in the protocol stack. Once the bypass is set up, the adapter is able to receive packets which are part of that flow specification. For example, in one example the packets are expected to be responses from a READ
15 RAW request for a block of data (block 201). The network adapter upon receiving the packet, and before passing the packet up the stack determines whether the packet falls within the set up that defines the flow specification for this read bypass session (block 202). If not, then the packet is processed normally (block 203) and the algorithm ends (block 204). If the packet falls
20 within the set up parameters at block 202, then the header of the packet is stored in a working buffer accessible by the driver, including the sequence number and length parameters from the header (block 205). Using this data, the protocol stack is capable of performing protocol maintenance functions necessary for maintaining the session. Also the payload is moved directly to the target buffer
25 (block 206). The payload move process may occur before, after or in parallel with storing the header in the working buffer. Next, if necessary for a given implementation, the set up data for the receive bypass is updated based on the information in the header (block 207). Finally the process ends (block 204).

5 The foregoing description of a preferred embodiment of the invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obviously, many modifications and variations will be apparent to practitioners skilled in this art. It is intended that the scope of the invention be defined by the following claims and their equivalents.

What is claimed is:

CLAIMS

1 1. A method for transferring data on a network from a data source
2 to an end station executing a multi-layer network protocol, including a network
3 layer and at least one higher layer, through a network interface on the end
4 station, comprising:

5 receiving in the network interface a packet which carries a data payload
6 from a block of data in the data source, and a control field identifying the
7 packet;

8 determining based on the control field in the network interface whether
9 the packet matches a flow specification, and if so transferring the data payload
10 in the packet directly to a target buffer assigned by a process at a layer higher
11 than the network layer.

1 2. The method of claim 1, wherein the control field in the packet
2 includes a packet header.

1 3. The method of claim 1, wherein the network protocol comprises
2 TCP/IP, and the packet control field comprises a TCP/IP header.

1 4. The method of claim 1, including prior to receiving the packet,
2 allocating the target buffer for the plurality of packets, and notifying the
3 network interface of the allocated target buffer.

1 5. The method of claim 1, the network interface is coupled to a
2 network medium supporting a maximum packet size, and including transmitting
3 a request from an application for transfer of a block of data from the data
4 source, the block of data having a length potentially greater than the maximum
5 packet size for the medium.

1 6. The method of claim 5, including notifying the network interface
2 in response to the request of a flow specification for the block of data according
3 to the multi-layer network protocol, and wherein the step of receiving the packet
4 includes identifying packet using the flow specification.

1 7. The method of claim 6, wherein the network protocol comprises
2 TCP/IP, and the flow specification includes a sequence number of a first byte
3 from the plurality of packets to be stored in the target buffer.

1 8. The method of claim 1, wherein the flow specification includes
2 a sequence number for the block of data.

1 9. The method of claim 8, wherein the flow specification includes
2 IP source and destination addresses and TCP port numbers.

1 10. A method for transferring data on a network from a data source
2 to an end station executing a multi-layer network protocol through a network
3 interface on the end station, including medium access control layer processes,
4 comprising:

5 establishing a connection with a destination for a session according to
6 the network protocol;

transmitting a request for transfer of a block of data from the data source, and providing a flow specification and an identifier of a target buffer to the network interface;

receiving in the network interface a plurality of packets which carry respective data payloads, packets in the plurality of packets including control fields identifying whether the packet falls within the flow specification of the block of data,

upon receiving a packet, determining in the network interface whether the packet falls within the flow specification, and if so transferring the data payload to the target buffer.

11. The method of claim 10, wherein the control field in the first packet includes a packet header.

12. The method of claim 10, wherein the network protocol comprises TCP/IP, and the packet control data comprises a TCP/IP header.

13. The method of claim 10, wherein the network protocol comprises TCP/IP, and the flow specification includes a sequence number of a first byte from the plurality of packets to be stored in the target buffer.

14. The method of claim 10, wherein the flow specification includes a sequence number for the block of data.

15. The method of claim 14, wherein the flow specification includes IP source and destination addresses and TCP port numbers.

1 16. A method for transferring data on a network from a data source
2 to an end station executing a TCP/IP network protocol through a network
3 interface on the end station including medium access control layer processes
4 below TCP/IP, comprising:

5 establishing a connection with a destination for a session according to
6 the TCP/IP network protocol;

7 transmitting a request from the application, for transfer of a block of data
8 from the data source, and providing a flow specification for the block of data
9 and an identifier of a target buffer to the network interface;

10 receiving in the network interface a plurality of packets which carry
11 respective data payloads from the block of data in the data source, and each
12 packet in the plurality of packets including a TCP/IP header,

13 upon receiving each packet, determining in the network interface
14 whether the packet falls within the flow specification, and if so
15 transferring the data payload to the target buffer.

1 17. The method of claim 16, wherein the flow specification includes
2 a sequence number for bytes of data in the block of data.

1 18. The method of claim 17, wherein the flow specification includes
2 IP source and destination addresses and TCP port numbers.

1 19. The method of claim 16, wherein the target buffer comprises a
2 buffer assigned at the TCP layer or higher.

1 20. The method of claim 16, wherein the target buffer comprises a
2 buffer assigned at a layer higher than the TCP layer.

RECEIVE PROCESSING WITH NETWORK PROTOCOL BYPASS

ABSTRACT

An adapter is provided with intelligence that allows it to separate the header parts of a packet being received from the payload it carries, and in most cases move the payload directly into a destination buffer at the application layer or file system layer. Copies by the intermediate layers of the protocol stack are bypassed, reducing the number of times that the payload of a communication must be copied by the host system. At the network interface, a plurality of packets is received, and the payload of each is bypassed directly into the target destination buffer. The network interface device identifies the packets which are in the sequence of packets carrying payload to be stored in the target buffer by the flow specification carried with such packets. Also, the packets carrying data payload for the file include a sequence number or other identifier by which the network interface is able to determine the offset within the target buffer to which the packet is to be stored.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

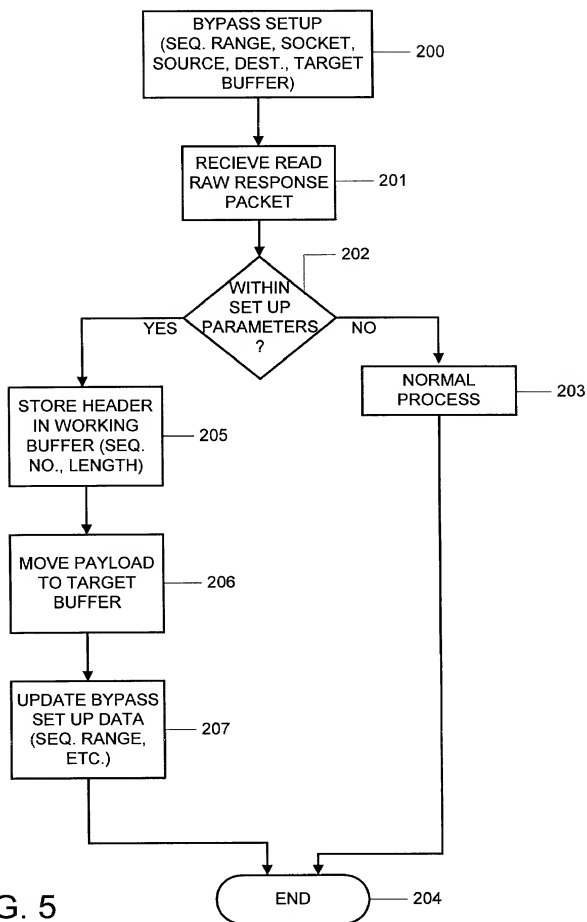


FIG. 5

**COMBINED DECLARATION AND POWER OF ATTORNEY
FOR UTILITY PATENT APPLICATION**

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

RECEIVE PROCESSING WITH NETWORK PROTOCOL BYPASS

the specification of which

 X is attached hereto.

_____ was filed on _____ as Application No. _____
and was amended on * _____
(If Applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56(a) which states in relevant part: "Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section....The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98."

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate as indicated below and have also identified below any foreign application for patent or inventor's certificate on this invention having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Claimed

_____	_____	_____	Yes	No
(Number)	(Country)	(Day/Month/Year Filed)		

_____	_____	_____	Yes	No
(Number)	(Country)	(Day/Month/Year Filed)		

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s), and under §119(e) of any United States provisional application(s), listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulation, §1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

_____	_____	_____
(Application Serial No.)	(Filing Date)	(Patented, Pending, Abandoned)

_____	_____	_____
(Application Serial No.)	(Filing Date)	(Patented, Pending, Abandoned)

I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith, and to file, prosecute and to transact all business in connection with international applications directed to said invention:

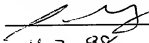
Paul Davis	29,294
Mark A. Haynes	30,846
John J. Bruckner	35,816
Charles C. Cary	36,764
Michael J. Murphy	37,404
David J. Weitz	38,362
Kent R. Richardson	39,443
David J. Abraham	39,554
U.P. Peter Eng	39,666
George A. Willman	41,378

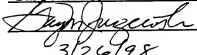
Address all correspondence to:

Mark A. Haynes
Wilson Sonsini Goodrich & Rosati
650 Page Mill Road
Palo Alto, CA 94304

Direct all telephone calls to Mark A. Haynes at (650) 493-9300.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Title 18, United States Code, §1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or
first inventor: Glenn William Connery
Inventor's signature: 
Date: 4-2-98
Citizenship: Canada
Residence: 655 S. Fair Oaks #B301, Sunnyvale, California 94086
Post Office Address: Same as above.

Full name of second joint
inventor, if any: Gary Jaszewski
Inventor's signature: 
Date: 3/26/98
Citizenship: United States
Residence: 14900 Quito Road, Los Gatos, California 95030
Post Office Address: Same as above.

Full name of third joint
inventor, if any:

Richard Reid

Inventor's signature:

Richard Reid

Date:

3-27-98

Citizenship:

United States

Residence:

465 Stierlin Road, Apt. 15, Mountain View, California 94043

Post Office Address:

Same as above.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION

IN RE PATENT APPLICATION OF)

Glenn William Connery,)
Gary Jaszewski, and)
Richard Reid)

Examiner: Not Yet Assigned

Application No. Not Yet Assigned)

Group Art Unit: Not Yet Assigned

Filing Date: Herewith)

Title: RECEIVE PROCESSING WITH)
NETWORK PROTOCOL BYPASS)

POWER OF ATTORNEY BY ASSIGNEE
TO EXCLUSION OF INVENTOR UNDER 37 C.F.R. § 3.71
WITH REVOCATION OF PRIOR POWERS

Commissioner of Patents
and Trademarks
Washington, D.C. 20231

Sir:

The undersigned ASSIGNEE of the entire interest in the above-identified application for letters patent hereby appoints:

Mark A. Haynes	30,846
Paul Davis	29,294
Kent R. Richardson	39,443
Charles C. Cary	36,764
George A. Willman	41,378
David J. Weitz	38,362
John J. Bruckner	35,816
U.P. Peter Eng	39,666
David Abraham	39,554

to prosecute this application and transact all business in the United States Patent and Trademark Office in connection therewith and hereby revokes all prior powers of attorney; said appointment to be to the exclusion of the inventors and the inventors' attorneys in accordance with the provisions of 37 C.F.R. § 3.71.

The following evidentiary documents establish a chain of title from the original owner to the Assignee:

X a copy of an Assignment attached hereto, which Assignment has been (or is herewith) forwarded to the Patent and Trademark Office for recording; or

___ the Assignment recorded on _____ at reel ___, frames ___ - ___.

Pursuant to 37 C.F.R. § 3.73(b) the undersigned Assignee hereby states that evidentiary documents have been reviewed and hereby certifies that, to the best of ASSIGNEE's knowledge and belief, title is in the identified ASSIGNEE.

Direct all telephone calls to Mark A. Haynes, (650) 493-9300.

Address all correspondence to:

Mark A. Haynes
WILSON SONSINI GOODRICH & ROSATI
650 Page Mill Road
Palo Alto, California 94304-1050

ASSIGNEE: 3Com Corporation

Name: Mark Michael
(Signature)

Name: Mark Michael
(Print or Type)

Title: Senior Vice President

Date: 4/2/98